arXiv:1611.05080v2 [q-bio.NC] 13 Jan 2017

Neural Stochastic Codes, Encoding and Decoding

Hugo Gabriel Eyherabide

Department of Computer Science and Helsinki Institute for Information Technology, University of Helsinki
Gustaf Hällströmin katu 2b, FI00560, Helsinki, Finland

Email: neuralinfo@eyherabidehg.com **Weppage:** eyherabidehg.com

Keywords: Neural codes, reduced codes, stochastic codes, neural encoding, neural decoding, observer perspective, spike-time precision, discrimination, noise correlations, information theory, mismatched decoding

Abstract

Understanding brain function, constructing computational models and engineering neural prosthetics require assessing two problems, namely encoding and decoding, but their relation remains controversial. For decades, the encoding problem has been shown to provide insight into the decoding problem, for example, by upper bounding the decoded information. However, here we show that this need not be the case when studying response aspects beyond noise correlations, and trace back the actual causes of this major departure from traditional views. To that end, we reformulate the encoding and decoding problems from the observer or organism perspective. In addition, we study the role of spike-time precision and response discrimination, among other response aspects, using stochastic transformations of the neural responses, here called stochastic codes. Our results show that stochastic codes may cause different information losses when used to describe neural responses and when employed to train optimal decoders. Therefore, we conclude that response aspects beyond noise correlations may play different roles in encoding and decoding. In practice, our results show for the first time that decoders constructed low-quality descriptions of response aspects may operate optimally on high-quality descriptions and vice versa, thereby potentially yielding experimental and computational savings, as well as new opportunities for simplifying the design of computational brain models and neural prosthetics.

1 Introduction

Assessing what aspects of neural activity are informative —the encoding problem— and what role they play in brain computations —the decoding problem— are two of the most fundamental questions in neuroscience, but their relation remain controversial. Traditionally, these two problems have been related through data processing theorems which ensure, for example, that the decoded information cannot exceed the encoded information (Schneidman et al., 2003; Cover and Thomas, 2006; Victor and Nirenberg, 2008; Eyherabide and Samengo, 2013). Although seemingly valid when studying specific response aspects such as first-spike latencies and spike counts (Furukawa and Middlebrooks, 2002; Nelken and Chechik, 2007; Gaudry and Reinagel, 2008; Eyherabide et al., 2008, 2009; Eyherabide and Samengo, 2010a,b), this relation has recently been shown invalid when studying covariations between response aspects under repeated stimulation, also called noise correlations. Contrary to previously thought, here we show that this relation is also invalid when studying response aspects beyond noise correlations, and establish the actual conditions under which the relation holds.

Specifically, previous studies have shown that, although not always (Pillow et al., 2008; Eyherabide et al., 2009), decoding algorithms that ignore noise correlations may cause negligible information losses even in situations within which noise correlation have been found to increase the encoded information (Nirenberg et al., 2001; Latham and Nirenberg, 2005; Oizumi et al., 2009; Ince et al., 2010; Oizumi et al., 2010; Pita-Almenar et al., 2011; Meytlis et al., 2012; Eyherabide and Samengo, 2013; Eyherabide, 2016; Oizumi et al., 2016). This paradoxical observation has previously been attributed to at least four different causes: inherent limitations of the employed measures, putative fundamental differences between ignoring response aspects and response probabilities, neural activity that would only occur should neurons fire independently but not in the recorded data, or the fact that, unlike latencies or spike counts, noise correlations cannot be ignored using deterministic functions of the recorded neural responses, also called reduced codes (Nirenberg and Latham, 2003; Schneidman et al., 2003; Latham and Nirenberg, 2005; Averbeck et al., 2006). Although still subject to active research (Oizumi et al., 2009; Ince et al., 2010; Oizumi et al., 2010; Eyherabide and Samengo, 2013; Eyherabide, 2016), this paradox and its causes continue to be regarded as a sole hallmark of studying noise correlations.

This belief notwithstanding, here we show that analogous observations arise when studying other response aspects such as spike-time precision or response discrimination. This finding constitutes a huge departure from traditional view, because it arises even when the response aspects are defined through stimulus-independent transformations, seemingly violating the data processing theorems. To that end, we reinterpret

the encoding and the decoding problems from the observer perspective, and disentangle them based on their operation, as opposed to the type of measure used to assess them (Quiroga and Panzeri, 2009). Then, we study the role of spike-time precision and response discrimination in encoding and decoding using stochastic codes, as opposed to reduced codes (Victor and Purpura, 1996; Victor, 2005; Nelken, 2008; Quiroga and Panzeri, 2009; Kayser et al., 2010; Rusu and Florian, 2014). We show that, like noise correlations, stochastic codes may yield different information losses when employed within the encoding and the decoding problems, but for none of the causes previously proposed by noise-correlation studies (Nirenberg and Latham, 2003; Schneidman et al., 2003; Latham and Nirenberg, 2005; Averbeck et al., 2006). Finally, we determine under which conditions the encoding and decoding problems are reliably related, thereby clarifying traditional views. We conclude that response aspects including and beyond noise correlations may play different roles in encoding and decoding, and that decoders designed with noisy data may perform optimally when operating on quality data. In this way, we open up new opportunities for studying neural codes and brain computations that may potentially yield experimental and computational savings and reduce the complexity and cost of neural prosthetics.

2 Methods

2.1 Statistical notation

When no risk of ambiguity arises, we here employ the standard abbreviated notation of statistical inference (Casella and Berger, 2002), denoting random variables with letters in upper case, and their values, with the same letters but in lower case. Accordingly, the symbols $P(x|y)$ and $P(x)$ always denote the conditional and unconditional probabilities, respectively, of the random variable X taking the value x given that the random variable Y takes the value y. Although this notation is common in neuroscience (Nirenberg and Latham, 2003; Quiroga and Panzeri, 2009; Latham and Roudi, 2013), we have found that it may potentially lead to confusing the encoding and the decoding problems. In those cases, we employ the more general symbols $P(X{=}v|Y{=}w)$ and $P(X{=}v)$ to denote the conditional and unconditional probabilities, respectively, that the random variable X takes the value v given that the random variable Y takes the value w.

2.2 Encoding

The process of converting stimuli S into neural responses $\mathbf{R}$ (e.g., spike-trains, local-field potentials, electroencephalographic or other brain signals, etc.) is called encoding (Schneidman et al., 2003; Panzeri et al., 2010). The information that $\mathbf{R}$ contains about S is here quantified as

$$I(S;\mathbf{R}) = \sum_{s,\mathbf{r}} P(s,\mathbf{r}) \log \frac{P(s|\mathbf{r})}{P(s)} \,. \tag{1}$$

More generally, the information $I(S;X)$ that any random variable X, including but not limited to $\mathbf{R}$, contains about S can be computed using the above formula with $\mathbf{R}$ replaced by X. For compactness, we will denote $I(S;X)$ as I_X unless ambiguity arises.

2.3 Decoding

The process of transforming $\mathbf{R}$ into estimated stimuli $\hat{S}$ (or into perceptions, decisions and actions) is called decoding (Schneidman et al., 2003; Panzeri et al., 2010). Under the Bayesian coding hypothesis (Knill and Pouget, 2004; van Bergen et al., 2015), decoding is usually performed using optimal decoders (also called Bayesian or maximum-a-posteriori decoders, and ideal homunculus or observers, among other names; see Eyherabide (2016) and references therein). In their most general form, these decoders map $\mathbf{R}$ into $\hat{S}$ according to the following formula

$$\hat{s} = \arg\max_{s} P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r}) \,. \tag{2}$$

The symbol $\hat{\mathbf{R}}$ denotes the surrogate responses either used to train the optimal decoder, or associated with the approximation $P(s|\hat{\mathbf{r}})$ of the real posterior probabilities $P(s|\mathbf{r})$ used to construct the decoder (Merhav et al., 1994; Nirenberg and Latham, 2003; Latham and Nirenberg, 2005; Quiroga and Panzeri, 2009; Oizumi et al., 2010; Eyherabide and Samengo, 2013; Latham and Roudi, 2013; Eyherabide, 2016).

The joint probability $P(s,\hat{s})$ of the presented and estimated stimuli, also called confusion matrix (Quiroga and Panzeri, 2009), can be computed as follows

$$P(s,\hat{s}){=}\sum_{s,\mathbf{r}_{\hat{s}}} P(s,\mathbf{r}) \,, \tag{3}$$

where $\mathbf{r}_{\hat{s}}$ denotes all $\mathbf{r}$ that are mapped into $\hat{s}$. Then, the information that $\hat{S}$ preserves about S is given by $I_{\hat{S}}$, whereas the decoding accuracy above chance level is here defined as follows

$$A_{\mathbf{R}}^{\hat{\mathbf{R}}} = \sum_s P(S{=}s, \hat{S}{=}s) - \max_s P(s)\,, \tag{4}$$

with the subscript and superscript in A_X^Y always indicating the neural codes used to construct and to operate optimal decoders, in that order. The data processing theorems ensure that both $A_{\mathbf{R}}^{\mathbf{R}} \geq A_{\hat{\mathbf{R}}}^{\hat{\mathbf{R}}}$ and $I_{\mathbf{R}} \geq I_{\hat{S}}$.

2.4 Neural codes

Any representation of the recorded neural activity into a vector $\mathbf{R}{=}[R_1, \ldots, R_N]$ of N response aspects is usually called a neural code. The aspects may represent, for example, neurons, neural populations, or cortical areas (Nirenberg and Latham, 2003; Schneidman et al., 2003; Latham and Nirenberg, 2005; Eyherabide and Samengo, 2013; Eyherabide, 2016), but also first-spike latencies, spike counts, and spike-timing variability (Furukawa and Middlebrooks, 2002; Nelken and Chechik, 2007; Eyherabide and Samengo, 2010b; Eyherabide, 2016). To assess what aspects are informative, $\mathbf{R}$ is usually transformed into another neural code $\hat{\mathbf{R}}$, often called reduced code (Schneidman et al., 2003; Eyherabide and Samengo, 2010b), through stimulus-independent deterministic functions that typically preserve a limited number of response aspects or a coarser version of the existing ones. As we note here, this assessment can also be conducted using stimulus-independent stochastic functions, here called stochastic codes (see Results). In both cases, the data processing theorem ensures that $I_{\mathbf{R}} \geq I_{\hat{\mathbf{R}}}$.

2.5 Noise correlations

Given a neural code $\mathbf{R}{=}[R_1, \ldots, R_N]$, the aspects are said noise independent if $P(\mathbf{r}|s){=}P^{NI}(\mathbf{r}|s)$ for all $\mathbf{r}$ and s, with $P^{NI}(\mathbf{r}|s){=}\prod_{n=1}^{N} P(r_n|s)$. Otherwise, they are said noise correlated. As mentioned in Eyherabide (2016), this definition should not be confused with both those that average over stimuli (Gawne and Richmond, 1993; Pereda et al., 2005; Womelsdorf et al., 2012), and therefore prone to cancellation effects and to confusing noise correlations with activity correlations (Nirenberg and Latham, 2003; Schneidman et al., 2003; Eyherabide, 2016); and those that are limited to specific linear or nonlinear types of correlations (Pereda et al., 2005; Cohen and Kohn, 2011; Latham and Roudi, 2013).

2.6 Encoding-oriented measures

From an information-theoretical standpoint, the importance of response aspects in neural encoding has previously been quantified in many ways (see Eyherabide and Samengo, 2010b, and references therein), out of which we chose the following four measures

$$\Delta I_{\hat{\mathbf{R}}} = I_{\mathbf{R}} - I_{\hat{\mathbf{R}}} \tag{5}$$

$$\Delta A_{\hat{\mathbf{R}}} = A_{\mathbf{R}}^{\mathbf{R}} - A_{\hat{\mathbf{R}}}^{\hat{\mathbf{R}}} \tag{6}$$

$$\Delta I_{\hat{\mathbf{S}}} = I_{\mathbf{R}} - I_{\hat{\mathbf{S}}} \tag{7}$$

$$\Delta I_{\hat{S}} = I_{\mathbf{R}} - I_{\hat{S}} \tag{8}$$

Here, $\mathbf{R}$ and $\hat{\mathbf{R}}$ are two arbitrary neural codes; whereas $\hat{\mathbf{S}}$ and $\hat{S}$ denote a sorted stimulus list and the most-likely stimulus, both according to $P(S{=}s|\hat{\mathbf{R}}{=}\hat{\mathbf{r}})$. The measure $\Delta I_{\hat{\mathbf{S}}}$ is here introduced as an encoding-oriented version of the decoding-oriented measure ΔI^{LS} previously proposed by Ince et al. (2010) and defined after the next two sections together with other decoding-oriented measures. Notice that $\Delta I_{\hat{\mathbf{S}}}{=}\Delta I_{\hat{S}}$ when the number of stimuli is two.

In this study, stochastic functions always exist that transform $\mathbf{R}$ into $\hat{\mathbf{R}}$ in a stimulus-independent manner, unless otherwise stated. Hence, the transformation from $\mathbf{R}$ into $\hat{S}$ can be interpreted as a sequential process (Geisler, 1989; Eyherabide and Samengo, 2013) in which $\mathbf{R}$ is first transformed into $\hat{\mathbf{R}}$, then into $\hat{\mathbf{S}}$, and finally into $\hat{S}$. Consequently, $\Delta I_{\hat{\mathbf{R}}}$, $\Delta I_{\hat{\mathbf{S}}}$ and $\Delta I_{\hat{S}}$ can be interpreted as accumulated information losses after the first, second and third transformations, respectively, whereas $\Delta A_{\hat{\mathbf{R}}}$, as the accuracy loss after the first transformation. The data processing theorems ensure that all the above measures are never negative.

The above four measures are here regarded as encoding-oriented, even though the last three use optimal decoders. Indeed, our classification does not follow previous criteria based on the nature of the measure (Schneidman et al., 2003; Quiroga and Panzeri, 2009), which are often obscure and arguably questionable (Eyherabide, 2016). Specifically, consider $\Delta I_{\hat{\mathbf{R}}}$ and $\Delta I_{\hat{S}}$. The former has always been regarded as an encoding measure even when no stimulus-independent function from $\mathbf{R}$ into $\hat{\mathbf{R}}$ exists (Nirenberg and Latham, 2003), and even though it need not be conclusive about the actual encoding mechanism (see Results). Classifying the latter as an encoding measure may seem questionable because it is based on encoded information, but in the output of a decoder (Quiroga and Panzeri, 2009; Eyherabide, 2016). However, in the light of previous studies on the role of noise correlations in neural decoding (Nirenberg et al., 2001; Nirenberg and Latham, 2003; Latham and Nirenberg, 2005; Eyherabide, 2016), classifying it as a decoding measure may also seem questionable because the decoder is trained and tested with the same responses.

2.7 Noise correlations in neural encoding

Notice that Eq (5) has previously been used to quantify the importance of noise correlations in encoding information (Schneidman et al., 2003; Eyherabide et al., 2009). In those cases, $\hat{\mathbf{R}}$ has been replaced with surrogate responses generated assuming that neurons are noise independent, here denoted $\mathbf{R}^{NI}$. However, we can prove the following theorem:

Theorem 1. $\mathbf{R}^{NI}$ *may not be related to* $\mathbf{R}$ *through stochastic codes and is never related to* $\mathbf{R}$ *through a reduced code.*

Proof. The first part is proved in Results. The second part was first proved in Schneidman et al. (2003), but their proof is invalid when $\mathbf{R}^{NI}$ contains the same responses as $\mathbf{R}$. In that case, we can prove the second part by contradiction, assuming that a mapping exists from $\mathbf{R}$ into $\mathbf{R}^{NI}$ which is deterministic, and hence bijective. Therefore, both $\mathbf{R}^{NI}$ and $\mathbf{R}$ maximize the conditional entropy given S over the probability distributions with the same the marginals. Because the probability distribution achieving this maximum is unique Cover and Thomas (2006), $\mathbf{R}^{NI}$ and $\mathbf{R}$ must be the same, thereby proving our statement. ■

Consequently, $\Delta I_{\mathbf{R}^{NI}}$, $\Delta A_{\mathbf{R}^{NI}}$, $\Delta I_{\mathbf{S}^{NI}}$, and $\Delta I_{S^{NI}}$ (with $\mathbf{S}^{NI}$ and S^{NI} analogous to $\hat{\mathbf{S}}$ and $\hat{S}$, respectively, but for optimal decoders operating on $\mathbf{R}^{NI}$) can be ensured neither to be nonnegative nor to compare information of the same type (Nirenberg and Latham, 2003), unless a stochastic code exists that maps $\mathbf{R}$ into $\mathbf{R}^{NI}$.

2.8 Decoding-oriented measures

From and information-theoretical standpoint, the importance of response aspects in neural decoding has been previously quantified in many ways (see Eyherabide and Samengo, 2013, and references therein), out of which we chose the following five that are more closely related to optimal decoding

$$\Delta I^{D} = \sum_{s,\mathbf{r}} P(s,\mathbf{r}) \ln \frac{P(s|\mathbf{r})}{P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r})} \tag{9}$$

$$\Delta I^{DL} = \min_{\theta} \sum_{s,\mathbf{r}} P(s,\mathbf{r}) \ln \frac{P(s|\mathbf{r})}{P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r},\theta)} \tag{10}$$

$$\Delta I^{LS} = I_{\mathbf{R}} - I_{\hat{\mathbf{S}}} \tag{11}$$

$$\Delta I^{B} = I_{\mathbf{R}} - I_{\hat{S}} \tag{12}$$

Here, $\mathbf{R}$ is the neural code on which the decoders operate; $\hat{\mathbf{R}}$, the neural code with which the decoders are trained; θ, a real scalar; K, the number of stimuli; $\hat{\mathbf{S}}$ and $\hat{S}$, a

sorted stimulus list and the most-likely stimulus, both according to $P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r})$; and we define $P(s|\hat{\mathbf{r}},\theta)$ below, after resolving the flaws of previous definitions (Merhav et al., 1994; Latham and Nirenberg, 2005; Oizumi et al., 2009, 2010; Eyherabide and Samengo, 2013; Oizumi et al., 2016). Specifically, we can prove the following theorem

Theorem 2. *$P(s|\hat{\mathbf{r}},\theta){\propto}P(s)$ for all s and $\hat{\mathbf{r}}$ if a $\check{\mathbf{r}}$ exists such that $P(\mathbf{R}{=}\check{\mathbf{r}}|S{=}s){>}P(\hat{\mathbf{R}}{=}\check{\mathbf{r}}|s){=}0$ for some or all s.*

Proof. Using the standard conventions that $0\log 0{=}0$ and $x\log 0{=}\infty$ for $x{>}0$ (Cover and Thomas, 2006), Eq (B13a) in Latham and Nirenberg (2005) is fulfilled for all $P(s|\hat{\mathbf{r}},\theta)$ such that $P(s|\hat{\mathbf{r}},\theta){>}0$ when $P(\mathbf{R}{=}\hat{\mathbf{r}}|S{=}Vs){>}P(\hat{\mathbf{r}}|s){=}0$. Our result immediately follows by solving Eq (B15) in Latham and Nirenberg (2005) with $\beta{=}0$. ■

In addition, we can generalize the flaws of previous definitions of $P(s|\hat{\mathbf{r}},\theta)$ found in Eyherabide and Samengo (2013) as follows

Theorem 3. *$P(s|\hat{\mathbf{r}},\theta){<}\infty$ if $P(\hat{\mathbf{r}}|s){=}P(\mathbf{R}{=}\hat{\mathbf{r}}){=}0$ regardless of s. Otherwise, $P(s|\hat{\mathbf{r}},\theta){=}0$ if $P(\hat{\mathbf{r}}|s){=}P(\mathbf{R}{=}\hat{\mathbf{r}}|S{=}s){=}0$ for some but not all s.*

Proof. Using the standard convention that $0\log 0{=}0$, Eq (B13a) in Latham and Nirenberg (2005) is fulfilled only if $P(s,\hat{\mathbf{r}}|\theta){=}0$ when $P(s,\hat{\mathbf{r}}){=}0$ and $P(s,\mathbf{r}){=}0$ for $\mathbf{r}{=}\hat{\mathbf{r}}$. Our result immediately follows using Bayes' rule. ■

Consequently, here we define $P(s|\hat{\mathbf{r}},\theta)$ as follows

$$P(s|\hat{\mathbf{r}},\theta) \propto \begin{cases} P(s) & \text{if } \exists s,\check{\mathbf{r}} \text{ such that } P(\mathbf{R}{=}\check{\mathbf{r}}|S{=}s){>}P(\hat{\mathbf{R}}{=}\check{\mathbf{r}}|\mathbf{R}{=}s){=}0 \\ 0 & \text{if } P(\hat{\mathbf{r}}|s){=}P(\mathbf{R}{=}\hat{\mathbf{r}}|S{=}s){=}0 \text{ for some but not all } s \\ P(s)\,P(\hat{\mathbf{r}}|s)^{\theta} & \text{otherwise} \end{cases} \tag{13}$$

In addition, we will compute the accuracy loss from the decoding perspective as follows

$$\Delta A^{B} = A^{\mathbf{R}}_{\mathbf{R}} - A^{\hat{\mathbf{R}}}_{\mathbf{R}}\,. \tag{14}$$

This measure differs from $\Delta A_{\hat{\mathbf{R}}}$ in that the last term uses optimal decoders that operate on $\mathbf{R}$, as opposed to $\hat{\mathbf{R}}$. Notice that the decoding-oriented measures defined above have previously been used to quantify the importance of noise correlations in neural decoding, for example, by replacing $\hat{\mathbf{R}}$ with $\mathbf{R}^{\mathit{NI}}$ (Nirenberg et al., 2001; Nirenberg and Latham, 2003; Latham and Nirenberg, 2005; Eyherabide et al., 2009; Oizumi et al., 2009; Quiroga and Panzeri, 2009; Ince et al., 2010; Oizumi et al., 2010; Eyherabide and Samengo, 2013; Latham and Roudi, 2013; Eyherabide, 2016).

2.9 Shortcomings of decoding-oriented measures

The measures ΔI^{D}, ΔI^{DL}, ΔI^{LS} and ΔI^{B} have all been previously regarded as quantifying the information loss caused when optimal decoders operate on $\mathbf{R}$, but make decisions assuming that the input is $\hat{\mathbf{R}}$. The first measure, ΔI^{D}, was introduced by Nirenberg et al. (2001) with $P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r}){=}P^{NI}(s|\mathbf{r})$, and later extended to any $P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r})$ (Nirenberg and Latham, 2003; Quiroga and Panzeri, 2009; Latham and Roudi, 2013; Eyherabide, 2016). Although often overlooked (Quiroga and Panzeri, 2009; Latham and Roudi, 2013), ΔI^{D} presents the following two difficulties: first, it can exceed $I_{\mathbf{R}}$ and tend to infinity if $P(s|\hat{\mathbf{r}}){=}0$ when $P(s|\mathbf{r}){>}0$ for some s and $\mathbf{r}{=}\hat{\mathbf{r}}$ (Schneidman et al., 2003); and second, here we note that it may be undefined if $P(\hat{\mathbf{r}}){=}0$ when $P(\mathbf{r}){>}0$ for some $\mathbf{r}{=}\hat{\mathbf{r}}$ (Fig 3B).

The second measure, ΔI^{DL}, was introduced by Latham and Nirenberg (2005) and has always been deemed exact (Latham and Nirenberg, 2005; Oizumi et al., 2009; Quiroga and Panzeri, 2009; Ince et al., 2010; Oizumi et al., 2010; Latham and Roudi, 2013; Oizumi et al., 2016). Latham and Nirenberg (2005) showed that, unlike ΔI^{D}, $\Delta I^{DL}{\leq}I_{\mathbf{R}}$, but their proof ignored the cases mentioned in our theorems 2 and 3. Nevertheless, these theorems imply that $\Delta I^{DL}{\leq}I_{\mathbf{R}}$ through Eq (10). However, Eyherabide and Samengo (2013) has shown that ΔI^{DL} may exceed both ΔI^{LS} and ΔI^{B}; and Eyherabide (2016), that it may overestimate the loss when decoding together neural populations that transmit independent information.

The third measure, ΔI^{LS}, was introduced by Ince et al. (2010) and quantifies the difference between two encoded informations: the one in $\mathbf{R}$, and the one in the output of decoders that, after observing $\mathbf{r}$, produce a stimulus list sorted according to $P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r})$. The fourth measure, ΔI^{B}, quantifies the difference between the information encoded in $\mathbf{R}$ and that encoded in the output of an optimal decoder constructed using $P(s|\hat{\mathbf{r}})$. When the number of stimuli is two, ΔI^{LS} reduces to ΔI^{B} (Eyherabide and Samengo, 2013; Eyherabide, 2016).

As we show in Fig 3B, ΔI^{D}, ΔI^{LS}, ΔI^{B}, and ΔA^{B} are undefined if the actual responses $\mathbf{R}$ are not contained within the surrogate responses $\hat{\mathbf{R}}$. The result for ΔI^{D} was already explained in the first paragraph. For the other measures, our observation stems from the fact that, for the missing responses, an optimal decoder simply does not know what output should be produced. Our observation does not apply to ΔI^{DL} due to our theorems 2 and 3, but can be easily seen from the definition of ΔI^{DL} as a communication information loss (Latham and Nirenberg, 2005; Oizumi et al., 2010; Eyherabide, 2016). Indeed, the fact that optimal decoders may fail to operate on $\mathbf{R}$ due to the missing responses in $\hat{\mathbf{R}}$ indicates that no information can be reliably transmitted, thereby yielding $\Delta I^{DL}{=}I_{\mathbf{R}}$ (Latham and Nirenberg, 2005; Eyherabide, 2016).

Finally, notice that decoding measures may yield different and sometimes contradic-

tory results (Oizumi et al., 2009, 2010; Eyherabide and Samengo, 2013; Latham and Roudi, 2013). However, the contradictions can be partially resolved by noticing that they have been previously derived using different notions of information. The reader is referred to Eyherabide (2016) and references therein for further remarks on these measures and notions.

3 Results and Discussion

3.1 Encoding and decoding from the observer perspective

The role of first-spike latencies, spike counts, or other response aspects in encoding information — the encoding problem — and in brain computations — the decoding problem — has previously been assessed, for example, using methods based either on information theory or on decoders (Quiroga and Panzeri, 2009). For decades, methods based on information theory have often been related to the encoding problem, whereas methods based on decoders have often been related to the decoding problem. However, here we argue that the nature of the employed method is not reliably related to its operational significance, at least from the observer perspective.

Information-based methods typically compare the information encoded in two representations of the neural response, usually called neural codes and here denoted $\mathbf{R}$ and $\hat{\mathbf{R}}$. The code $\mathbf{R}$ typically preserves all response aspects, whereas the code $\hat{\mathbf{R}}$ usually arises from some transformation of $\mathbf{R}$ that ignores some response aspects, or at least the additional information they carry. The ignored response aspects are usually deemed inessential for encoding information when the information in $\mathbf{R}$ and $\hat{\mathbf{R}}$ is approximately equal, and important otherwise. However, the encoded information need not provide conclusive insight about the encoding mechanisms, but about the read out mechanisms.

To illustrate this, consider the hypothetical experiment shown in Fig 1A. There, the responses of a single neuron elicited by two visual stimuli have been characterized by counting the number of spikes within consecutive time-bins starting from the stimulus onset. This characterization preserves all the encoded information regardless of whether the bin size is 5, 10, or $15\,ms$, thereby shedding limited insight into the actual spike-time precision employed by the encoding mechanism, which is $10\,ms$. However, this result certainly proves that such knowledge is inessential for ideal observers, who can read the spike trains without loss using any of the aforementioned bin sizes.

Decoder-based methods usually compare the decoded informations or the decoding accuracies of two decoders: one constructed using $\mathbf{R}$ and another one constructed using $\hat{\mathbf{R}}$. The ignored response aspects are usually deemed inessential for decoding

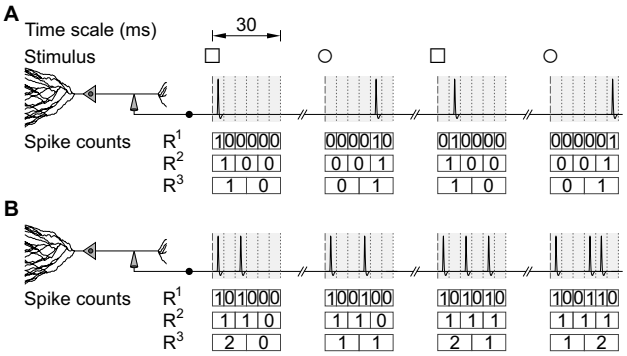

Figure 1: **Lossless codes need not disentangle the encoding mechanisms.** A: Hypothetical intracellular recording of the precise spike patterns elicited by a single neuron after presenting in alternation two visual stimuli, namely □ and ◯. Stimulus-response probabilities are arbitrary, and time is discretized in bins of $5\,ms$. The responses are recorded within $30\,ms$ time-windows after stimulus onset. Each type of frame elicits single spikes fired with different latencies, uniformly distributed between 0 and $10\,ms$ after the onset of □, and between 20 and $30\,ms$ after the onset of ◯. Responses have been characterized by counting the number of spikes within consecutive time-bins of size 5, 10 and $15\,ms$ starting from the stimulus onset, thereby yielding a discrete-time sequences here denoted $\mathbf{R}^1$, $\mathbf{R}^2$ and $\mathbf{R}^3$, respectively. B: Analogous description to panel A, but with each type of frame producing two different types of response patterns composed of 2 or 3 spikes.

information when the information extracted by these two decoders is approximately equal, and important otherwise. However, the information extracted by optimal decoders trained and tested with surrogate responses $\mathbf{R}^{NI}$ generated assuming that neurons are noise independent (Quiroga and Panzeri, 2009) need not provide insight into the necessity for taking into account noise correlations when constructing optimal decoders (Nirenberg and Latham, 2003; Eyherabide, 2016).

Our observations motivate us to reformulate the coding problem from the perspective of an ideal observer or organism (Fig 2A). Within this perspective, we devise the encoding problem as concerning the losses caused when the ideal observers either can only see and are only allowed to see not the actual neural responses ($\mathbf{R}$), but a transformed version of them ($\hat{\mathbf{R}}$) that preserves a limited number of response aspects (Fig 2C). The decoding problem is here interpreted analogously to the decoding perspective of the role of noise correlations, namely, as concerning the losses caused when observers make decisions assuming that they see $\hat{\mathbf{R}}$ instead of $\mathbf{R}$ (Fig 2D). Mathematically, both problems can be tackled by training ideal observers not with $\mathbf{R}$ but with $\hat{\mathbf{R}}$ (Fig 2B). However, within the encoding problem, the ideal observer operates on $\hat{\mathbf{R}}$, whereas within the decoding problem, the ideal observer operates on $\mathbf{R}$.

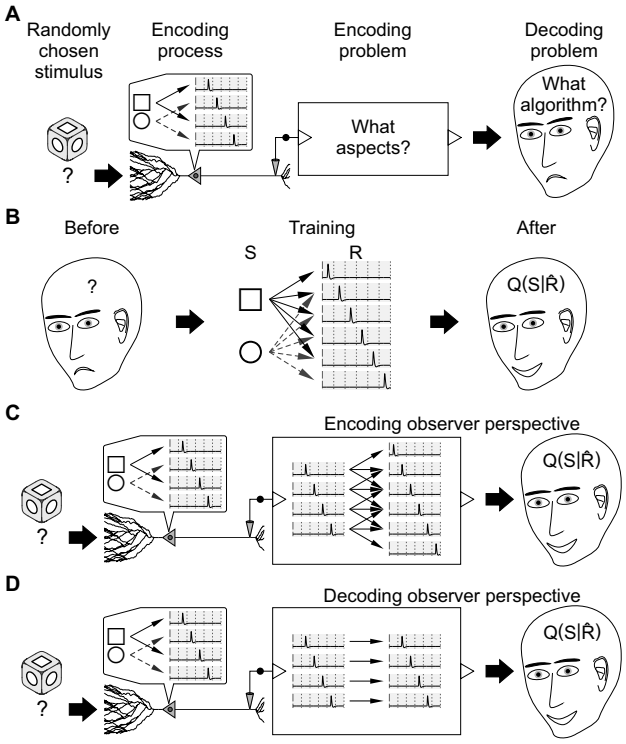

Figure 2: **Encoding and decoding problems from the organism perspective.** A: The encoding problem is here reinterpreted as concerning the losses caused by transforming the recorded neural responses $\mathbf{R}$ into another responses $\hat{\mathbf{R}}$ before supplying them to ideal observers (Nelken and Chechik, 2007); whereas the decoding problem, as concerning the losses caused by training ideal observers with transformed responses $\hat{\mathbf{R}}$ or approximate response likelihoods, as opposed to the actual responses that the observer sees (Nirenberg and Latham, 2003; Latham and Nirenberg, 2005; Quiroga and Panzeri, 2009; Latham and Roudi, 2013). B: Both problems can be tackled by training ideal observers with the transformed responses $\hat{\mathbf{R}}$. C: However, in the encoding problem, the above observer is provided with the neural responses after the transformation, namely $\hat{\mathbf{R}}$. D: On the contrary, in the decoding problem, the above observer is provided with the original neural responses unchanged, namely $\mathbf{R}$.

3.2 Relation between encoding and decoding

These conceptual and operational differences notwithstanding, decoding-oriented measures of loss have previously been regarded as generalizing their encoding-oriented counterparts (Nirenberg et al., 2001; Nirenberg and Latham, 2003; Quiroga and Panzeri, 2009; Latham and Roudi, 2013). Unfortunately, these claims ought to be observed with caution for at least three reasons: they need not be accurate, potentially confuse encoding and decoding, and are limited to specific types of codes. In this section, we will explain and justify these reasons.

Concerning the accuracy of the claims, Nirenberg and Latham (2003) claimed that their version ΔI^D for studying noise correlations is always positive and reduces to the difference in encoded information between two neural codes: one using large time-bins and another one using small time-bins. However, this difference can be positive or negative, regardless of the order in which the informations are subtracted, thereby contradicting the putative positivity of both ΔI^D and the measure they proposed. Indeed, in Fig 1B, using time-bins of 5 or $15\,ms$ preserves all the encoded information, but using intermediate time-bins of $10\,ms$ preserves none, thereby proving our statement.

Our finding seemingly contradicts previous experimental studies (Reinagel and Reid, 2000; Quiroga and Panzeri, 2013) and, most importantly, the data processing theorem (Cover and Thomas, 2006). Indeed, it shows that previous experimental findings of encoded information monotonically decreasing with bin size are not completely attributable to time discretization. However, it does not violate the data processing theorem because the code with $15\,ms$ time-bins cannot be derived from that with $10\,ms$ time-bins through stimulus-independent transformations.

Actually, the study of Nirenberg and Latham (2003) proposed a generalization of ΔI^D, as opposed to a reduction of it, but unfortunately it cannot be interpreted in the same way. Indeed, their generalization compares the additional costs of using two different codes for constructing two decoders, respectively, that are then employed to decode a third code. Thus, this generalization need not be suitable for addressing the decoding problem, which concerns the losses caused by decoding a code assuming that it is another one, unless it is first reduced to ΔI^D.

The confusion between encoding and decoding arises from the particular choice of probabilities that previous studies (Nirenberg and Latham, 2003; Quiroga and Panzeri, 2009; Latham and Roudi, 2013) have employed when computing ΔI^D for stimulus-independent transformations of the original codes representing the neural responses, but we defer the proofs until the last section of Results. Most importantly, the above generalization and computations of ΔI^D, as conducted in previous studies (Nirenberg and Latham, 2003; Quiroga and Panzeri, 2009; Latham and Roudi, 2013), together with their criticism (Schneidman et al., 2003), are actually limited to stimulus-independent deterministic transformations, often called reduced codes, as opposed stochastic ones, here called stochastic codes. In the next section we characterize this codes, and in the one immediately after, we will assess what role they play in encoding and decoding from the observer perspective.

3.3 Stochastic codes generalize reduced codes

Stochastic codes arise naturally when studying the importance of specific response aspects (Eyherabide and Samengo, 2010b). However, stochastic codes are most valuable when studying other response aspects such as spike-time precision (Kayser et al., 2010; Quiroga and Panzeri, 2009), response discrimination (Victor and Purpura, 1996; Victor, 2005; Rusu and Florian, 2014), noise along neural pathways (Nelken, 2008), and even noise correlations. In these cases, the use of reduced codes may be questionable, if not impossible. In this section, we will first illustrate these claims before giving a formal definition of stochastic codes.

Consider the hypothetical experiment in Fig 3A, in which the neural responses $\mathbf{R}=[L, C]$ can be completely characterized by the first-spike latencies (L) and the spike counts (C). The importance of C can be studied not only using a reduced code that replaces all its values with a constant (Fig 3B), but also a stochastic code that shuffles the values of C across all responses with the same L (Fig 3C), or even a stochastic code that preserves the original value of L but chooses the value of C from some possibly $L-$dependent arbitrary probability distribution (Fig 3D).

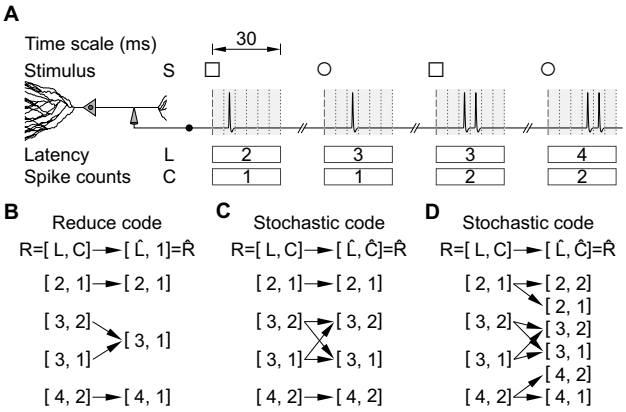

Figure 3: **Stochastic codes include and generalize reduced codes.** A: Analogous description to Fig 1A, but with responses characterized using a code $\mathbf{R}=[L, C]$ that preserves the first-spike latency (L) and the spike-count (C). B: Deterministic transformation (arrows) of $\mathbf{R}$ in panel A into a reduced code $\hat{\mathbf{R}}=[L, 1]$, which ignores the additional information carried in C by considering it constant and equal to unity. This reduced code can also be reinterpreted as a stochastic code with transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ defined by Eq (3.4). C: The additional information carried in C is here ignored by shuffling the values of C across all trails with the same L, thereby turning $\mathbf{R}$ in panel A into a stochastic code $\hat{\mathbf{R}}=[\hat{L}, \hat{C}]$ with transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ defined by Eq (3.4). D: The additional information carried in C is here ignored by replacing the actual value of C for one chosen with some possibly L-dependent probability distribution (Eq (3.4)).

Using either reduced or stochastic codes in the previous example seemingly yields the same result (but see next section). However, this need not always be the case. For example, spike-time precision has previously been studied using at least three different methods, here called binning, bin-shuffling and jittering. Binning is perhaps the most common of the three methods and consists in transforming the recorded spike trains into discrete-time sequences (Fig 1), thereby yielding a reduced code that can be interpreted as a stochastic code (Fig 3). Bin-shuffling consists in randomly permuting spikes and silences within consecutive non-overlapping windows (Fig 4A), which inherently constitutes a stochastic code.

These two methods are equivalent in the sense that the responses generated by each method can be related through stimulus-independent stochastic functions to those generated by the other, and therefore contain the same amount and type of information. However, both methods suffer from the same drawback: they both treat spikes differently depending on their location within the time window. Indeed, both methods preserve the distinction between two spikes located at different time bins, but not within the same time bin, even if the separation between the spikes is the same.

The third method, jittering, consists in shuffling the recorded spikes within time windows centered at each spike (Fig 4B). Notice that the responses generated by this method need not be related to the responses generated by the other two methods through stimulus-independent stochastic functions, nor vice versa. However, this method inherently yields a stochastic code, and, unlike the previous methods, treats all spikes in the same manner.

As another example, consider the effect of response discrimination, as studied in the seminal work of Victor and Purpura (1996). There, two responses were considered indistinguishable when some measure of distance between the responses was less than some predefined value. However, the distances were there used for transforming neural responses through a method based on cross-validation that, as we note here, is not guaranteed to be stimulus-independent as stochastic codes. Other methods exist which simply merge indistinguishable responses thereby yielding a reduced code, but these methods are limited to distances that are transitive.

Here we overcome these limitations by devising a method based on stochastic codes, which mimics the operation of an observer that confuses responses by randomly treating them as if it was another response. This method includes those based on reduced codes, and even those based on possibly non-symmetrical and non-metric measures of response similarity. Unlike the previous methods, this one can effectively be applied to the example shown in Fig 4C.

Formally, Stochastic codes are here defined as neural codes $\hat{\mathbf{R}}$ that can be obtained through stimulus-independent stochastic functions of other codes $\mathbf{R}$. After observing

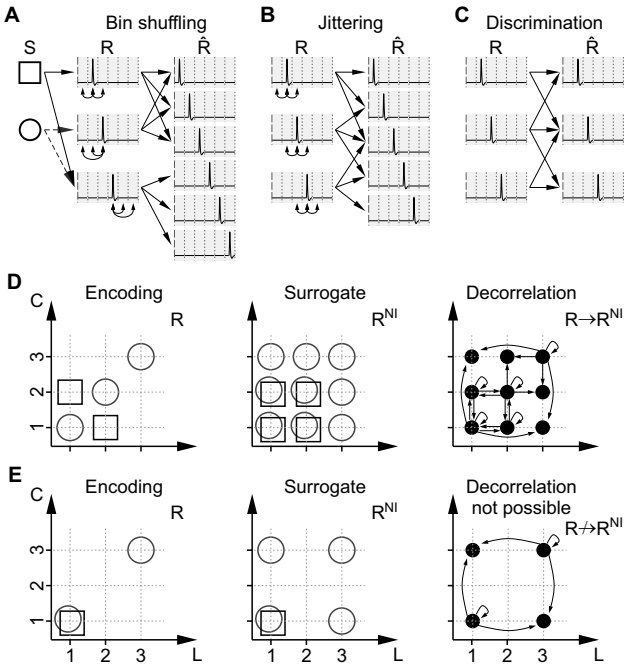

Figure 4: **Examples of stochastic codes.** Hypothetical experiments analogous to that described in Fig 3A. A: Stochastic function (arrows on the left) modeling the encoding process, followed by another one with transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ given by Eq (3.4). This function turns $\mathbf{R}$ into a stochastic code $\hat{\mathbf{R}}$ by shuffling spikes and silences within consecutive bins of $15\,ms$ starting from stimulus onset. Notice that $\hat{\mathbf{R}}$ is analogous to $\mathbf{R}^3$ in Fig 1 (see Results). B: Responses $\mathbf{R}$ in panel A, followed by a stochastic function with $Q(\hat{\mathbf{R}}|\mathbf{R})$ given by Eq (3.4), which introduces jitter uniformly distributed within $15\,ms$ windows centered at each spike. C: Responses $\mathbf{R}$ in panel A, followed by a stochastic function with $Q(\hat{\mathbf{R}}|\mathbf{R})$ given by Eq (3.4), which models the inability to distinguish responses with spikes occurring in adjacent bins, or equivalently, with distances $D^{spike}[q{=}1]{\leq}1$ or $D^{interval}[q{=}1]{\leq}1$ (see Victor and Purpura (1996) and Victor (2005) for further remarks on these distances). Notice that $\hat{\mathbf{R}}$ contains the same responses as $\mathbf{R}$. D: Cartesian coordinates depicting: on the left, responses $\mathbf{R}$ of a neuron different from that in panel A, with L and C are positively correlated when elicited by ◯, and negatively correlated when elicited by □; in the middle, the surrogate responses $\mathbf{R}^{NI}$ that would occur should L and C be noise independent (middle); and on the right, a stimulus-independent stochastic function that turns $\mathbf{R}$ into $\mathbf{R}^{NI}$ with $Q(\mathbf{R}^{NI}|\mathbf{R})$ given by Eq (3.4). E: Same description as in panel D, but with L and C noise independent given □, and with the stochastic function depicted on the right turning $\mathbf{R}$ into $\mathbf{R}^{NI}$ given ◯ but not □.

that $\mathbf{R}$ adopted the value $\mathbf{r}$, these functions produce a single value $\hat{\mathbf{r}}$ for $\hat{\mathbf{R}}$ chosen with transition probabilities $Q(\hat{\mathbf{r}}|\mathbf{r})$ such that

$$P(\hat{\mathbf{r}}|s) = \sum_{\mathbf{r}} P(\mathbf{r}|s)\, Q(\hat{\mathbf{r}}|\mathbf{r})\,. \tag{15}$$

When for each $\mathbf{r}$ exists a $\hat{\mathbf{r}}$ such that $Q(\hat{\mathbf{r}}|\mathbf{r})=1$, stochastic codes become reduced codes (Fig 3).

The problem of finding a feasible $Q(\hat{\mathbf{r}}|\mathbf{r})$ can be readily solved using linear programming (Boyd and Vandenberghe, 2004), which we can simplify by noticing the following two properties

1) $Q(\hat{\mathbf{r}}|\mathbf{r})=0$ if $P(s,\mathbf{r})>0$ and $P(s,\hat{\mathbf{r}})=0$, for otherwise it would contradict the definition of stochastic codes; and

2) A feasible $Q(\hat{\mathbf{r}}|\mathbf{r})$ never exists if $I_{\hat{\mathbf{R}}}>I_{\mathbf{R}}$, for otherwise it would contradict the data processing theorem.

For example, in Fig 4D, we can decorrelate first-spike latencies (L) and spike counts (C) by modeling $\mathbf{R}^{NI}$ as a stochastic code derived from $\mathbf{R}$. To that end, we can solve Eq (15) for $Q(\hat{\mathbf{R}}|\mathbf{R})$ but with $\hat{\mathbf{R}}$ replaced by $\mathbf{R}^{NI}$. In principle, the condition $I_{\mathbf{R}}>I_{\mathbf{R}^{NI}}$ always hold, and thus a feasible $Q(\mathbf{R}^{NI}|\mathbf{R})$ may exist (property 2). Such probability must be zero at least whenever $\mathbf{R}^{NI}\in\{[1,3];[2,3];[3,3];[3,2];[3,1]\}$ and $\mathbf{R}\in\{[1,2];[2,1]\}$ (property 1). One possible solution is given by Eq (3.4).

However, notice that stochastic codes need not always exist. For example, in Fig 4E, the condition $I_{\mathbf{R}}<I_{\mathbf{R}^{NI}}$ always hold, and therefore no stochastic code can map $\mathbf{R}$ into $\mathbf{R}^{NI}$. Notice that Schneidman et al. (2003) employed an analogous example, but involving different neurons instead of response aspects, only for comparing ΔI^{D} with $\Delta I_{\hat{\mathbf{R}}}$, and ignoring whether $\mathbf{R}^{NI}$ constitutes as a stochastic code. Finally, stochastic codes must not be confused with probabilistic population codes (Knill and Pouget, 2004), stimulus-dependent noise, decoding algorithms based on cross-validation (Victor and Purpura, 1996; Quiroga and Panzeri, 2009), and list decoding (Schneidman et al., 2003; Ince et al., 2010; Eyherabide and Samengo, 2013).

3.4 The role of stochastic codes

Like reduced codes, stochastic codes are here defined as stimulus-independent transformations of another code. Accordingly, both types of codes preserve only information contained in the original code that limits the decoded information when decoders operate on them. Although stochastic codes may include more responses than the original codes, this should arguably raise no concerns (Nirenberg and Latham, 2003; Schneidman et al., 2003) because the transformations that define them are stimulus independent and can be implemented during decoding. Thus, for stochastic codes, the encoding and the decoding problems may be related as they are for reduced codes, even when arising from ignoring noise correlations (Fig 4D).

We tested this hypothesis by comparing three encoding-oriented measures of information loss ($\Delta I_{\hat{\mathbf{R}}}$, $\Delta I_{\hat{\mathbf{S}}}$, $\Delta I_{\hat{S}}$) and one of accuracy loss ($\Delta A_{\hat{\mathbf{R}}}$), with four decoding-oriented measures of information loss (ΔI^{D}, ΔI^{DL}, ΔI^{LS}, ΔI^{B}) and one of accuracy loss (ΔA^{B}). Their definitions, rationale, and classification are given in Methods. We found that, for the stochastic codes in Figures 3–5, the encoding-oriented measures were greater or less than the decoding-oriented measures depending on the case and the probabilities (Table 1). Consequently, our results refute the above hypothesis and, most importantly, prove that previous controversies about encoding and decoding transcend the study of noise correlations.

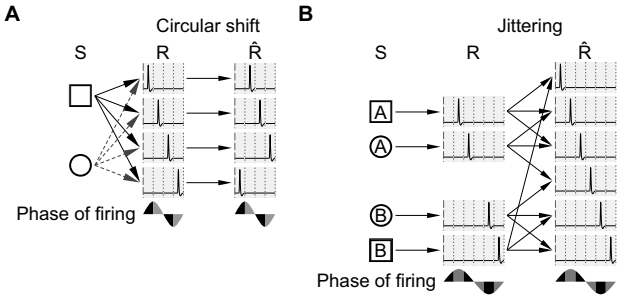

Figure 5: **Stochastic codes may play different roles in encoding and decoding.** A: Hypothetical experiment with two visual stimuli, namely □ and ◯. The stimuli are transformed (left arrows) into neural responses containing a single spike ($C{=}1$) fired at different phases (Φ) with respect to a cycle of $20\,ms$ period starting at stimulus onset. The phases have been discretized in intervals of size $\pi/2$ and wrapped to the interval $[0, 2\pi)$; the resulting value is also known as principal value. The encoding process is followed by a circular phase-shift that transforms $\mathbf{R}{=}\Phi$ into another code $\hat{\mathbf{R}}{=}\hat{\Phi}$ with transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ defined by Eq (3.4). Notice that $\hat{\mathbf{R}}$ contains all and only responses present in $\mathbf{R}$. B: Analogous description to panel A, except that stimuli are four (🄰, Ⓐ, 🄱, and Ⓑ), whereas phases are measured with respect to a cycle of $30\,ms$ period and discretized in intervals of size $\pi/3$. The encoding process is followed by a stochastic transformation (right arrows) that introduces jitter, thereby transforming $\mathbf{R}{=}\Phi$ into another code $\hat{\mathbf{R}}{=}\hat{\Phi}$ with transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ defined by Eq (3.4).

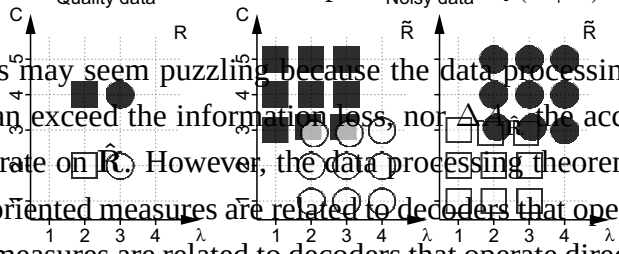

Our conclusions may seem puzzling because the data processing theorems ensure that neither $\Delta I_{\hat{\mathbf{R}}}$ can exceed the information loss, nor $\Delta A_{\hat{\mathbf{R}}}$ the accuracy loss, caused when decoders operate on $\hat{\mathbf{R}}$. However, the data processing theorems are not violated because encoding-oriented measures are related to decoders that operate on $\hat{\mathbf{R}}$, whereas decoding-oriented measures are related to decoders that operate directly on $\mathbf{R}$ (Fig 2D).

To gain additional insight, we divided the stochastic codes in four groups. In the first group (Fig 4D), noise correlations can be ignored through stochastic codes, but the results remain unchanged. Consequently, contrary to previously thought (Nirenberg and Latham, 2003), the discrepancy between encoding- and decoding-oriented measures cannot be

Table 1: **The encoding and decoding problems are unrelated for response aspects beyond noise correlations.** Maximum and minimum differences between encoding-oriented measures of information and accuracy losses and their decoding-oriented counterparts. For each example, the values were computed through the function `fminsearch` of Matlab 2016, with random initial values for the stimulus-response probabilities and the transition matrices. The computation was repeated until 20 consecutive runs failed to improve the estimate. The values are expressed as percentages of $I_{\mathbf{R}}$ (the information encoded in $\mathbf{R}$) or $A^{\mathbf{R}}_{\mathbf{R}}$ (the maximum accuracy above chance level when decoders operate on $\mathbf{R}$). Notice that $\Delta I_{\hat{\mathbf{S}}}{=}\Delta I_{\hat{S}}$ and $\Delta I^{LS}{=}\Delta I^{B}$ because all the cases comprise only two stimuli, and that the absolute value of $\Delta A_{\hat{\mathbf{R}}}{-}\Delta A^{B}$ can become extremely large when $A^{\mathbf{R}}_{\mathbf{R}}{\approx}0$.

Cases		Fig 4D	Fig 4B	Fig 4C	Fig 3D	Fig 4A	Fig 3B	Fig 5A
$\Delta I_{\hat{\mathbf{R}}}{-}\Delta I^{D}$	min	-79	-51	-40	0	0	—	-1193
	max	26	32	46	0	0	—	-171
$\Delta I_{\hat{\mathbf{R}}}{-}\Delta I^{DL}$	min	-34	-32	-18	0	0	-100	-100
	max	59	41	77	0	0	0	-100
$\Delta I_{\hat{\mathbf{R}}}{-}\Delta I^{B}$	min	-67	-62	-57	-63	-87	—	-100
	max	57	81	89	0	0	—	-4
$\Delta I_{\hat{S}}{-}\Delta I^{D}$	min	-79	-48	-38	0	0	—	-1192
	max	74	92	92	63	87	—	-71
$\Delta I_{\hat{S}}{-}\Delta I^{DL}$	min	-34	-27	-18	0	0	-100	-89
	max	91	92	92	63	87	0	83
$\Delta I_{\hat{S}}{-}\Delta I^{B}$	min	-51	-31	-18	0	0	—	-95
	max	59	91	93	0	0	—	79
$\Delta A_{\hat{\mathbf{R}}}{-}\Delta A^{B}$	min	-384	-200	-167	0	0	—	∞
	max	95	67	100	0	0	—	0

completely attributed to surrogate responses generated through stimulus-dependent transformations, or containing responses with information not present in the original neural responses. However, we can prove the following theorem

Theorem 4. *The transition probabilities $Q(\mathbf{R}^{NI}|\mathbf{R})$ of stochastic codes that ignore noise correlations may depend both on the marginal likelihoods, and on the noise correlations.*

Proof. We can prove the dependency on the marginal likelihoods by computing $Q(\mathbf{R}^{NI}|\mathbf{R})$ for the hypothetical experiment of Fig 4D. After some straightforward but tedious algebra, Eq (15) yields, for the $\mathbf{R}^{NI}$ associated with $\square$, the following equation must hold

$$P(L{=}1|S{=}\square) = 0.5\,\left[\delta q + \left(\delta q^2 + 4\,Q([1,2]|[2,1])\right)^{0.5}\right]\,, \tag{16}$$

where $\delta q{=}Q([1,2]|[1,2]) - Q([1,2]|[2,1])$. Hence, any change in $P(L{=}1|S{=}\square)$ must be followed by some change in $Q(\mathbf{R}^{NI}|\mathbf{R})$, thereby proving that the latter cannot be independent of the former.

We can prove the dependency on the noise correlations by computing $Q(\mathbf{R}^{NI}|\mathbf{R})$ for the hypothetical experiment of Fig 6. After some straightforward but tedious algebra analogous to that in Fig 4D, Eq (15) yields that for the responses associated with $\square$, $Q(\mathbf{R}^{NI}|\mathbf{R})$ must have the following form

$$Q(\mathbf{R}^{NI}|\mathbf{R}) = \begin{bmatrix} x_1 + x_4 & x_2 + x_5 & x_3 + x_6 & 1 - \sum_{i=1}^{6} x_i \\ x_1 & x_2 & x_3 & 1 - \sum_{i=1}^{3} x_i \\ x_4 & x_5 & x_6 & 1 - \sum_{i=4}^{6} x_i \\ 0 & 0 & 0 & 1 \end{bmatrix}, \qquad (17)$$

or else it will depend on the noise correlations. Here, $\mathbf{R}{=}[l, c]$ is associated with the row of index $i{=}2\,l{-}2{+}c$; whereas $\mathbf{R}^{NI}{=}[\hat{l}, \hat{c}]$, with the column of index $j{=}2\,\hat{l}{-}2{+}\hat{c}$; and $x_1, \ldots, x_6$ are probabilities. To resolve for $x_1, \ldots, x_6$, consider the set of all response distributions with the same marginals as $\mathbf{R}^{NI}$ that can be turned into $\mathbf{R}^{NI}$ through stimulus-independent stochastic functions regardless of their noise correlations. This set includes $\mathbf{R}^{NI}$, and therefore, $Q(\mathbf{R}^{NI}|\mathbf{R})$ should be able to transform $\mathbf{R}^{NI}$ into itself. However, it turns out that this is only possible when $P(\mathbf{R}^{NI}|\square)$ is unity for $\mathbf{R}^{NI}{=}[2, 2]$ and zero otherwise, thereby proving that $Q(\mathbf{R}^{NI}|\mathbf{R})$ depends on the noise correlations in $\mathbf{R}$, even if the marginal probabilities are fixed. ■

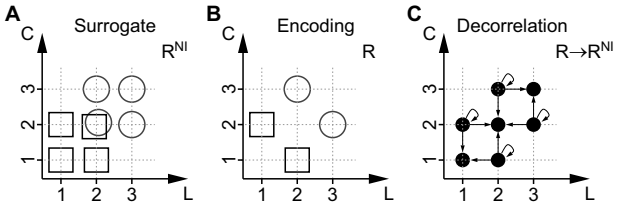

Figure 6: **Stochastic codes that ignore noise correlations may depend on them.** A: Cartesian coordinates representing a hypothetical experiment in which two different stimuli, namely $\square$ and $\bigcirc$, elicit single neuron responses ($\mathbf{R}^{NI}$) that completely characterized by their first-spike latency (L) and spike counts (C). Both L and C are noise independent. B: Cartesian coordinates representing a hypothetical experiment with the same marginal probabilities $P(l|s)$ and $P(c|s)$ as in A: depicting one among many possible types of noise correlations between L and C. C: Stimulus-independent stochastic function transforming the noise-correlated responses $\mathbf{R}$ of **b** into the noise-independent responses $\mathbf{R}^{NI}$ of **a**. The corresponding transition probability $Q(\mathbf{R}^{NI}|\mathbf{R})$ is given by Eq (3.4). Even when $P(l|s)$ and $P(c|s)$ are fixed, the transition probabilities may depend on the noise correlations, as proved in Results.

We tested if the discrepancy between encoding- and decoding-oriented measures stems from the aforementioned dependencies by comparing the values they attain within

the second group of stochastic codes (Fig 4 B,C). For these codes, which arise naturally when studying spike-time precision or response discrimination, we found that the dependencies never occur (namely, $Q(\hat{\mathbf{R}}|\mathbf{R})$ is independent on $P(S,\mathbf{R})$) but the results remain unchanged. Therefore, we conclude that the aforementioned discrepancy cannot be attributed to these dependencies, and can arise not only when studying noise correlations, but when studying other response aspects even if defined through constant and stimulus-response-independent transformations.

Further inspection shows that the discrepancy occurs even when stochastic codes only contain responses present in the original codes (Fig 4C), and both in the amount and in the type of information (Fig 5B). To prove the latter, consider that an ideal observer is trained using the noisy data $\hat{\mathbf{R}}$ shown in Fig 7A, but it is asked to operate on the quality data $\mathbf{R}$ shown in Fig 7B. The information losses $\Delta I_{\hat{\mathbf{R}}}$, ΔI^{D}, and ΔI^{DL} produced by this ideal observer are all equal to $I_{\mathbf{R}}/2$ regardless of whether $\mathbf{R}$ is transformed into $\hat{\mathbf{R}}$ before showing them to the ideal observer (the encoding problem) or not (the decoding problem).

However, the aforementioned transformation causes some responses to occur for all stimuli, thereby preserving only some information about frames ($I(S_F,\hat{\mathbf{R}})=I(S_F,\mathbf{R})/3$) and letters ($I(S_L,\hat{\mathbf{R}})=2\,I(S_L,\mathbf{R})/3$). In other words, within the encoding problem, $\hat{\mathbf{R}}$ causes a partial information loss $\Delta I_{\hat{\mathbf{R}}}$ that is composed of both frame and letter information. On the contrary, within the decoding problem, $\hat{\mathbf{R}}$ causes no information loss about letters but a total information loss about frames. We can prove the former by noticing that, for the responses that actually occur in $\mathbf{R}$, the ideal observer trained with $\hat{\mathbf{R}}$ can perfectly identify the letters (because $P(\hat{C}{=}2|S_L{=}\mathrm{A}){=}P(\hat{C}{=}4|S_L{=}\mathrm{B}){=}1$). We can prove the latter by noticing that $P(\hat{l}|\square){=}P(\hat{l}|\bigcirc)$ whenever $\hat{l}$ adopts a value that actually occurs in $\mathbf{R}$, namely 2 or 3. Analogous computations yield analogous results for the hypothetical experiment shown in Fig 5B.

These results contrast with those found in studies of noise correlations in two ways. First, those studies employed surrogate responses generated through stimulus-dependent transformations and often contain responses with information not present in the original neural responses. Second, for those surrogate responses, we can prove that ΔI^{DL} cannot reach $100\,\%$ unless $\Delta I_{\hat{\mathbf{R}}}$ does as well. Specifically, we can prove the following theorem

Theorem 5. *When ignoring noise correlations, $\Delta I^{DL}{=}I_{\mathbf{R}}$ if and only if $\Delta I_{\mathbf{R}^{NI}}{=}I_{\mathbf{R}}$, regardless of whether stochastic codes exist that map the actual responses $\mathbf{R}$ into the surrogate responses $\mathbf{R}^{NI}$ generated assuming noise independence.*

Proof. Consider a neural code $\mathbf{R}{=}[R_1,\ldots,R_N]$ and recall that the range of $\mathbf{R}^{NI}$ includes that of $\mathbf{R}$. Therefore, $\Delta I^{DL}{=}I_{\mathbf{R}}$ implies that the minimum in Eq (10) is attained when $\theta{=}0$. In that case, equation (B13a) in Latham and Nirenberg (2005) yields the

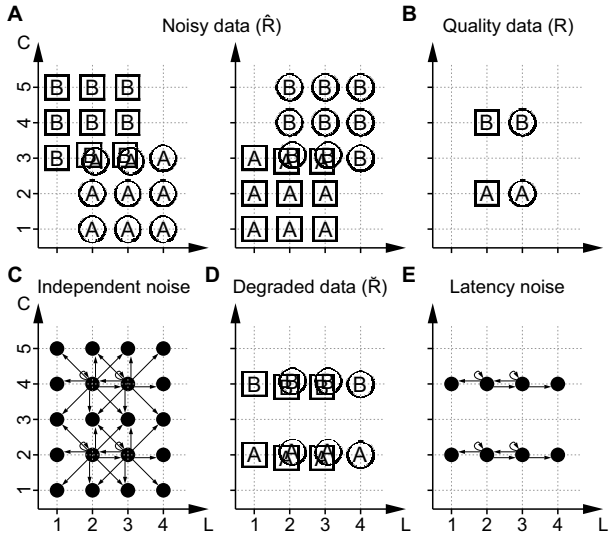

Figure 7: **Improving mismatched decoding by adding noise.** A: Noisy data recorded in a hypothetical experiment with four compound stimuli $S{=}[S_F, S_L]$, generated by choosing independently a frame ($S_F{=}\square$ or $\bigcirc$) and a letter ($S_L{=}\mathrm{A}$ or B), thereby yielding $\boxed{\mathrm{A}}$, $\textcircled{\mathrm{A}}$, $\boxed{\mathrm{B}}$, and $\textcircled{\mathrm{B}}$. Stimuli are transformed (left arrows) into neural responses $\hat{\mathbf{R}}{=}[\hat{L}, \hat{C}]$ with different number of spikes ($1{\leq}C{\leq}4$) fired at different first-spike latencies ($1{\leq}L{\leq}5$; time has been discretized in $5\,ms$ bins). Latencies are only sensitive to frames whereas spikes counts are only sensitive to letters, thereby constituting independent-information streams (namely $P(s, \mathbf{r}){=}P(s_F, l)\,P(s_L, c)$; Eyherabide, 2016). B: Quality data ($\mathbf{R}{=}[L, C]$) recorded as in panel A, but without noise. C: Stimulus-independent stochastic transformation with transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ given by Eq (3.4), that introduces independent noise both in the latencies and in the spike counts, thereby transforming $\mathbf{R}$ into $\hat{\mathbf{R}}$ and rendering $\hat{\mathbf{R}}$ as a stochastic code. D: Degraded data obtained by adding latency noise to the quality data. E: Representation of the stimulus-independent stochastic transformation with transition probabilities $Q(\breve{\mathbf{R}}|\mathbf{R})$ given by Eq (3.4) that adds latency noise in panel D.

following $\sum_{s,r_n} P(s, r_n)\,\log P(r_n|s){=}\sum_{s,r_n} P(s)\,P(r_n)\,\log P(r_n|s)$ for all $1{\leq}n{\leq}N$ for all n. After some more algebra and recalling that the Kullback-Leibler divergence is never negative, this equation becomes $I_{R_n}{=}0$. Consequently $\Delta I_{\mathbf{R}^{NI}}{=}I_{\mathbf{R}}$, thereby proving the "only if" part. For the "if" part, it is sufficient to notice that the last equality implies that $P^{NI}(\mathbf{r}|s){=}P^{NI}(\mathbf{r})$. ■

Consequently, our results cannot be inferred from the results obtained in studies of noise correlations.

In the third group (Fig 3D and Fig 4A), $\Delta I_{\hat{\mathbf{R}}}{=}\Delta I^{D}{=}\Delta I^{DL}$, $\Delta I_{\hat{\mathbf{S}}}{=}\Delta I^{LS}$, $\Delta I_{\hat{S}}{=}\Delta I^{B}$, and $\Delta A_{\hat{\mathbf{R}}}{=}\Delta A^{B}$. We can prove that these relations arise whenever the mapping from $\mathbf{R}$ into $\hat{\mathbf{R}}$ can be described using positive-diagonal idempotent stochastic matrices (Högnäs and Mukherjea, 2011). Specifically, we can prove the following theorem

Theorem 6. *Consider stimulus-independent stochastic functions f from a neural code $\mathbf{R}$ into another code $\hat{\mathbf{R}}$ which range $\mathcal{R}$ includes that of $\mathbf{R}$, and which transition probabilities $P(\hat{\mathbf{r}}|\mathbf{r})$ can be written as positive-diagonal idempotent right stochastic matrices with row and column indexes equally enumerating $\mathcal{R}$. Then, $\Delta I_{\hat{\mathbf{R}}}{=}\Delta I_{\breve{\mathbf{R}}}{=}\Delta I^{D}{=}\Delta I^{DL}$, $\Delta I_{\hat{\mathbf{S}}}{=}\Delta I^{LS}$, $\Delta I_{\hat{S}}{=}\Delta I^{B}$, and $\Delta A_{\hat{\mathbf{R}}}{=}\Delta A^{B}$, whenever reduced codes $\breve{\mathbf{R}}$ can be devised as lossless deterministic representations of $\hat{\mathbf{R}}$.*

Proof. The last condition assumes that a deterministic function g maps $\hat{\mathbf{R}}$ into $\breve{\mathbf{R}}$ such that $I_{\hat{\mathbf{R}}}{=}I_{\breve{\mathbf{R}}}$. Hence, $P(s|\hat{\mathbf{r}}){=}P(s|\breve{\mathbf{r}})$ when $\breve{\mathbf{r}}{=}g(\hat{\mathbf{r}})$, thereby carrying both $\hat{\mathbf{R}}$ and $\breve{\mathbf{R}}$ the same amount and type of information. However, the intention of this condition is more profound: It requires that, even when speaking of $\breve{\mathbf{R}}$, the decoding-oriented measures be applied to the stochastic code $\hat{\mathbf{R}}$. This avoids the problems found in Fig 3B, because the range $\mathcal{R}$ of $\hat{\mathbf{R}}$ includes that of $\mathbf{R}$. Futhermore, the restrictions on f imply that $\mathcal{R}$ can be partitioned into non-overlapping sets $\mathcal{R}_1, \ldots, \mathcal{R}_M$, each of which is mapped by f onto itself, and most importantly, that $P(\hat{\mathbf{r}}|\mathbf{r}){=}P(\hat{\mathbf{r}}|\mathcal{R}_m)$ when $\mathbf{r}{\in}\mathcal{R}_m$. Hence, for $\hat{\mathbf{r}}{\in}\mathcal{R}_m$, $P(\hat{\mathbf{r}}|s){=}P(\hat{\mathbf{r}}|\mathcal{R}_m)\,P(\mathcal{R}_m|s)$, thereby yielding $P(s|\hat{\mathbf{r}}){=}P(s|\mathcal{R}_m)$ and $P(s|\hat{\mathbf{r}},\theta){=}P(s|\mathcal{R}_m,\theta)$. Our result follow immediately after recalling that $P(s|\breve{\mathbf{r}}){=}P(s|\hat{\mathbf{r}})$. ■

This theorem is unnecessary for the quantitative equality between encoding- and decoding-oriented measures, but ensures that the equalities hold not only in amount but also in type. For example, in Fig 5B, we found that $\Delta I_{\hat{\mathbf{R}}}{=}\Delta I_{\hat{\mathbf{S}}}{=}\Delta I_{\hat{S}}{=}\Delta I^{D}{=}\Delta I^{DL}{=}\Delta I^{LS}{=}\Delta I^{B}{=}50\,\%$ and $\Delta A_{\hat{\mathbf{R}}}{=}\Delta A^{B}{\approx}67\,\%$ even though theorem 6 does not hold. However, the losses are not necessarily of the same type as we proved before. Most importantly, it formally justifies and clarifies the pervasive idea that, for reduced codes, the decoding-oriented measures reduce to their encoding-oriented counterparts. Indeed, the fourth group shows that, when not fulfilled, decoding-oriented measures can exceed their encoding-oriented counterparts (Fig 5A), or even be undefined (Fig 3B; see Methods). Consequently, contrary to previous studies (Nirenberg and Latham, 2003; Quiroga and Panzeri, 2009; Latham and Roudi, 2013), deterministic transformations need not ensure that the encoding and the decoding problems are related.

These results also reveal unexpected confounds in Nirenberg and Latham (2003); Quiroga and Panzeri (2009); Latham and Roudi (2013) that invalidate their conclusions by turning otherwise decoding-oriented measures into encoding-oriented ones. To reveal and resolve the flaws in the aforementioned studies, recall the experiment in Fig 5A. There, $\hat{\mathbf{R}}{=}f(\mathbf{R})$, with f being a deterministic bijective function. According to the above studies, ΔI^{D} should be computed through Eq (9) but with $P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r})$ replaced by $Q(s|\mathbf{r}){\propto}P(s)\,P(f(\mathbf{r})|s)$. Because f is bijective, $Q(s|\mathbf{r}){=}P(s|\mathbf{r}){=}P(S{=}s|\hat{\mathbf{R}}{=}f(\mathbf{r}))$. However, we can prove the following theorem

Theorem 7. *When deterministic functions f exist such that $\hat{\mathbf{R}}{=}f(\mathbf{R})$, replacing $P(S{=}s|\hat{\mathbf{R}}{=}\mathbf{r})$ with $P(S{=}s|\hat{\mathbf{R}}{=}f(\mathbf{R}))$ in Eq (9) turns ΔI^D into an encoding-oriented measure.*

Proof. Following the reasoning of previous studies (Nirenberg et al., 2001; Nirenberg and Latham, 2003; Latham and Nirenberg, 2005; Latham and Roudi, 2013), Eq (9) with the above replacement is analogous to comparing two optimal question-asking strategies induced by $P(s|\mathbf{r})$ and $P(s|\hat{\mathbf{r}})$, respectively. The former is applied directly on $\mathbf{R}$, but the latter is not. Instead, the latter is applied on $f(\mathbf{R}){=}\hat{\mathbf{R}}$. Consequently, the comparison addresses the effect of transforming $\mathbf{R}$ into $\hat{\mathbf{R}}$ before feeding them into the decoding process, thereby turning ΔI^D into an encoding-oriented measure. ■

Hence, our theorem proves that the computations in the aforementioned studies turn ΔI^D into an encoding-oriented measure, and that their conclusions confuse encoding and decoding. Consequently, caution must be exercise when computing the necessary probabilities for all decoding-oriented measures, including but not limited to ΔI^D.

In practice, our results open up the possibility of increasing the efficiency of decoders constructed with approximate descriptions of the neural responses, usually called approximate or mismatched decoders, by adding suitable amounts and types of noise to the decoder input. To see this, recall the example of Fig 7 in which a decoder constructed with noisy data ($\hat{\mathbf{R}}$; Fig 7A) was employed to decode quality data ($\mathbf{R}$; Fig 7B), thereby causing the information losses $\Delta I^D{=}\Delta I^{DL}{=}I_{\mathbf{R}}/2$. These losses can be decreased by feeding the decoder with a degraded version $\breve{\mathbf{R}}$ of the quality data (Fig 7D) generated through a stimulus-independent transformation that adds latency noise (Fig 7E). Decoding $\mathbf{R}$ as if it was $\hat{\mathbf{R}}$ by first transforming $\mathbf{R}$ into $\breve{\mathbf{R}}$ results in $\Delta I^D{=}\Delta I^{DL}{=}I_{\mathbf{R}}/3$, thereby recovering $33\,\%$ of the information previously lost. On the contrary, adding spike-count noise will tend to increase the losses. Thus, we have proved that adding suitable amounts and type of noise can increase the performance of approximate decoders even for response aspects beyond noise correlations. In addition, this result also indicates that, contrary to previously thought (Shamir, 2014), decoding algorithms need not match the encoding mechanisms for performing optimally from an information-theoretical standpoint.

Conclusion

Here we have reformulated the coding problem from the observer perspective as concerning two questions: what needs to be seen — the encoding problem — and how it must be interpreted — the decoding problem. These two problems were here shown to provide limited insight into each other when studying spike-time precision, response discrimination, or other aspects of the neural response. Furthermore, we have shown

that response aspects may play different roles in encoding and decoding even when defined through stimulus-independent transformations, and that decoding need not match encoding for response aspects beyond noise correlations. These findings constitute a major departure from traditional views on the role of response aspects in neural coding and brain computations. On the practical side, our most outstanding finding is that the decoded information need not be limited by the information carried in the data used to construct the decoder. This finding was here shown to open up new possibilities for increasing the performance of existing approximate decoders. Most importantly, it also indicates that decoders need not be retrained for operating optimally on data of higher quality than the one used in their construction, thereby potentially saving experimental and computational resources, and reducing the complexity and cost of neural prosthetics.

Acknowledgments

This work was supported by the Ella and Georg Ehrnrooth Foundation.

Appendix

Transition probabilities used in the figures

In Fig 3 , the transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ are defined as follows

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{for Fig 3B;} \tag{18}$$

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & a & \bar{a} & 0 \\ 0 & a & \bar{a} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad \text{for Fig 3C; and} \tag{19}$$

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \begin{bmatrix} b & \bar{b} & 0 & 0 & 0 & 0 \\ 0 & 0 & C & \bar{c} & 0 & 0 \\ 0 & 0 & c & \bar{c} & 0 & 0 \\ 0 & 0 & 0 & 0 & d & \bar{d} \end{bmatrix} \quad \text{for Fig 3D.} \tag{20}$$

Here, we have used the following conventions: row and column indexes enumerate $\mathbf{R}$ and $\hat{\mathbf{R}}$, respectively, in lexicographical order; $a{=}P(\mathbf{R}{=}[3,1])/P(L{=}3)$; $0{<}b,c,d{<}1$; and $\bar{x}{=}1{-}x$ for any number x. Notice that, the matrix in Eq (3.4) is positive-diagonal idempotent right stochastic as in theorem 6, whereas the one in Eq (3.4) can be made positive-diagonal idempotent by adding rows.

In Fig 4 , the transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ are defined using matrices in which row and column indexes enumerate $\mathbf{R}$ and $\hat{\mathbf{R}}$, respectively, according to their latencies in increasing order, thereby yielding

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \frac{1}{3}\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix} \text{ for Fig 4A;} \tag{21}$$

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \frac{1}{6}\begin{bmatrix} 2 & 2 & 2 & 0 & 0 \\ 0 & 2 & 2 & 2 & 0 \\ 0 & 0 & 2 & 2 & 2 \end{bmatrix} \text{ for Fig 4B;} \tag{22}$$

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \frac{1}{6}\begin{bmatrix} 3 & 3 & 0 \\ 2 & 2 & 2 \\ 0 & 3 & 3 \end{bmatrix} \text{ for Fig 4C; and} \tag{23}$$

$$Q(\mathbf{R}^{NI}|\mathbf{R}) = \frac{1}{2}\begin{bmatrix} 2b & \bar{b}c & \bar{c}\bar{b} & \bar{b}c & 0 & 0 & \bar{c}\bar{b} & 0 & 0 \\ \bar{a} & 2a & 0 & 0 & \bar{a} & 0 & 0 & 0 & 0 \\ a & 0 & 0 & 2\bar{a} & a & 0 & 0 & 0 & 0 \\ 0 & b & 0 & b & 2\bar{b}c & \bar{c}\bar{b} & 0 & \bar{c}\bar{b} & 0 \\ 0 & 0 & b & 0 & 0 & \bar{b}c & b & \bar{b}c & 2\bar{c}\bar{b} \end{bmatrix} \text{ for Fig 4D.} \tag{24}$$

Here, $a{=}P(\mathbf{R}{=}[1,2]|S{=}\square)$; $b{=}P(\mathbf{R}{=}[1,1]|S{=}\bigcirc)$; and $c{=}P(\mathbf{R}{=}[2,2]|S{=}\bigcirc)/\bar{b}$.

In Fig 5 , the transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ are defined using matrices in which row and column indexes indicate to the values of Φ and $\hat{\Phi}$, respectively, thereby yielding

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix} \text{ for Fig 5A; and} \tag{25}$$

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \frac{1}{3}\begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} \text{ for Fig 5B.} \tag{26}$$

In Fig 6 , the transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ can be obtained by solving Eq (15) and is given by the following

$$Q(\mathbf{R}^{NI}|\mathbf{R}) = 0.5\begin{bmatrix} \bar{a} & 2a & 0 & \bar{a} & 0 & 0 & 0 \\ a & 0 & 2\bar{a} & a & 0 & 0 & 0 \\ 0 & 0 & 0 & \underline{b} & 2\bar{b} & 0 & \underline{b} \\ 0 & 0 & 0 & \bar{b} & 0 & 2b & \bar{b} \end{bmatrix}, \tag{27}$$

with row and column indexes enumerating $\mathbf{R}$ and $\hat{\mathbf{R}}$, respectively, in lexicographical order; $a{=}P(\mathbf{R}{=}[1,2]|S{=}\square)$; and $b{=}P(\mathbf{R}{=}[3,2]|S{=}\bigcirc)$.

In Fig 7 , the transition probabilities $Q(\hat{\mathbf{R}}|\mathbf{R})$ can be defined using matrices in which row and column indexes enumerate $\mathbf{R}$ and $\hat{\mathbf{R}}$, respectively, in lexicographical order, thereby yielding

$$Q(\hat{\mathbf{R}}|\mathbf{R}) = \frac{1}{9}\begin{bmatrix} 1&1&1&0&1&1&1&0&1&1&1&0&0&0&0&0&0&0&0&0 \\ 0&1&1&1&0&1&1&1&0&1&1&1&0&0&0&0&0&0&0&0 \\ 0&0&0&0&0&0&0&0&1&1&1&0&1&1&1&0&1&1&1&0 \\ 0&0&0&0&0&0&0&0&0&1&1&1&0&1&1&1&0&1&1&1 \end{bmatrix} \text{ for Fig 7C; and} \tag{28}$$

$$Q(\breve{\mathbf{R}}|\mathbf{R}) = \frac{1}{3}\begin{bmatrix} 0&0&0&0&1&1&1&0&0&0&0&0&0&0&0&0&0&0&0&0 \\ 0&0&0&0&1&1&1&0&0&0&0&0&0&0&0&0&0&0&0&0 \\ 0&0&0&0&0&0&0&0&0&0&0&0&1&1&1&0&0&0&0&0 \\ 0&0&0&0&0&0&0&0&0&0&0&0&1&1&1&0&0&0&0&0 \end{bmatrix} \text{ for Fig 7E.} \tag{29}$$

References

Averbeck, B. B., Latham, P. E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.*, 7(5):358–366.

Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge university press.

Casella, G. and Berger, R. L. (2002). *Statistical Inference*. Duxbury, second edition edition.

Cohen, M. R. and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat Neurosci*, 14:811–819.

Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. Wiley-interscience, second edition edition.

Eyherabide, H. G. (2016. Preprint. Available from: arXiv:1608.05501. Cited 6 December 2016.). Disambiguating the role of noise correlations when decoding neural populations together.

Eyherabide, H. G., Rokem, A., Herz, A. V. M., and Samengo, I. (2008). Burst firing is a neural code in an insect auditory system. *Front. Comput. Neurosci.*, 2(3).

Eyherabide, H. G., Rokem, A., Herz, A. V. M., and Samengo, I. (2009). Bursts generate a non-reducible spike-pattern code. *Front. Neurosci.*, 3(1):8–14.

Eyherabide, H. G. and Samengo, I. (2010a). The information transmitted by spike patterns in single neurons. *J. Physiol. Paris*, 104(3):147–155.

Eyherabide, H. G. and Samengo, I. (2010b). Time and category information in pattern-based codes. *Front. Comput. Neurosci.*, 4(145).

Eyherabide, H. G. and Samengo, I. (2013). When and why noise correlations are important in neural decoding. *J. Neurosci.*, 33(45):17921–17936.

Furukawa, S. and Middlebrooks, J. C. (2002). Cortical representation of auditory space: information-bearing features of spike patterns. *J Neurophysiol*, 87(4):1749–1762.

Gaudry, K. S. and Reinagel, P. (2008). Information measure for analyzing specific spiking patterns and applications to lgn bursts. *Network*, 19:69–94.

Gawne, T. J. and Richmond, B. J. (1993). How independent are the messages carried by adjacent inferior temporal cortical neurons? *J. Neurosci.*, 13(7):2758–2771.

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychol. Rev.*, 96(2):267–314.

Högnäs, G. and Mukherjea, A. (2011). *Probability Measures on Semigroups: Convolution Products, Random Walks and Random Matrices*. Springer, New York, USA, second edition.

Ince, R. A. A., Senatore, R., Arabzadeh, E., Montani, F., Diamond, M. E., and Panzeri, S. (2010). Information-theoretic methods for studying population codes. *Neural Netw.*, 23(6):713–727.

Kayser, C., Logothetis, N. K., and Panzeri, S. (2010). Millisecond encoding precision of auditory cortex neurons. *Proc Natl Acad Sci*, 107(39):16976–16981.

Knill, D. C. and Pouget, A. (2004). The bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci*, 27(12):712–719.

Latham, P. E. and Nirenberg, S. (2005). Synergy, redundancy, and independence in population codes, revisited. *J. Neurosci.*, 25(21):5195–5206.

Latham, P. E. and Roudi, Y. (2013). *Role of correlations in population coding*, chapter 7, pages 121–138. CRC Press.

Merhav, N., Kaplan, G., Lapidoth, A., and Shamai Shitz, S. (1994). On information rates for mismatched decoders. *IEEE Trans Inf Theory*, 40(6):1953–1967.

Meytlis, M., Nichols, Z., and Nirenberg, S. (2012). Determining the role of correlated firing in large populations of neurons using white noise and natural scene stimuli. *Vision Res.*, 70:44–53.

Nelken, I. (2008). Processing of complex sounds in the auditory system. *Curr Opin Neurobiol*, 18(4):413–417.

Nelken, I. and Chechik, G. (2007). Information theory in auditory research. *Hear Res*, 229(1-2):94–105.

Nirenberg, S., Carcieri, S. M., Jacobs, A. L., and Latham, P. E. (2001). Retinal ganglion cells act largely as independent encoders. *Nature*, 411(6838):698–701.

Nirenberg, S. and Latham, P. E. (2003). Decoding neuronal spike trains: how important are correlations? *Proc. Natl. Acad. Sci. U.S.A.*, 100(12):7348–7353.

Oizumi, M., Amari, S. I., Yanagawa, T., Fujii, N., and Tsuchiya, N. (2016). Measuring integrated information from the decoding perspective. *PLoS Comput. Biol.*, 12(1):e1004654.

Oizumi, M., Ishii, T., Ishibashi, K., Hosoya, T., and Okada, M. (2009). A general framework for investigating how far the decoding process in the brain can be simplified. *Adv. Neural Inf. Process. Syst.*, pages 1225–1232.

Oizumi, M., Ishii, T., Ishibashi, K., Hosoya, T., and Okada, M. (2010). Mismatched decoding in the brain. *J. Neurosci.*, 30(13):4815–4826.

Panzeri, S., Brunel, N., Logothetis, N. K., and Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends Neurosci.*, 33(3):111–120.

Pereda, E., Quiroga, R. Q., and Bhattacharya, J. (2005). Nonlinear multivariate analysis of neurophysiological signals. *Prog. Neurobiol.*, 77(1–2):1–37.

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., and Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995–999.

Pita-Almenar, J. D., Ranganathan, G. N., and Koester, H. J. (2011). Impact of cortical plasticity on information signaled by populations of neurons in the cerebral cortex. *J Neurop*, 106:1118–1124.

Quiroga, R. Q. and Panzeri, S. (2009). Extracting information from neuronal populations: information theory and decoding approaches. *Nat. Rev. Neurosci.*, 10(3):173–185.

Quiroga, R. Q. and Panzeri, S. (2013). *Role of correlations in population coding*, chapter 8, pages 139–160. CRC Press.

Reinagel, P. and Reid, R. C. (2000). Temporal coding of visual information in the thalamus. *J Neurosci*, 20:5392–5400.

Rusu, C. V. and Florian, R. V. (2014). A new class of metrics for spike trains. *Neural Comput*, 26(2):306–348.

Schneidman, E., Bialek, W., and Berry, M. J. (2003). Synergy, redundancy, and independence in population codes. *J. Neurosci.*, 23(37):11539–11553.

Shamir, M. (2014). Emerging principles of population coding: in search for the neural code. *Curr Opin Neurobiol*, 25:140–148.

van Bergen, R. S., Ma, W. J., Pratte, M. S., and Jehee, J. F. M. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nat. Neurosci.*, 18(12):1728–1730.

Victor, J. D. (2005). Spike train metrics. *Curr Opin Neurobiol*, 15(5):585–592.

Victor, J. D. and Nirenberg, S. (2008). Indices for testing neural codes. *Neural comput*, 20:2895–2936.

Victor, J. D. and Purpura, K. P. (1996). Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J Neurophysiol*, 76(2):1310–1326.

Womelsdorf, T., Lima, B., Vinck, M., Oostenveld, R., Singer, W., Neuenschwander, S., and Fries, P. (2012). Orientation selectivity and noise correlation in awake monkey area v1 are modulated by the gamma cycle. *Proc Natl Acad Sci USA*, 109(11):4302–4307.